

EVOLUTION OF THE VELOCITY-DISPERSION FUNCTION OF LUMINOUS RED GALAXIES: A HIERARCHICAL BAYESIAN MEASUREMENT

YIPING SHU¹, ADAM S. BOLTON¹, DAVID J. SCHLEGEL², KYLE S. DAWSON¹, DAVID A. WAKE³, JOEL R. BROWNSTEIN¹,
 JON BRINKMANN⁴, AND BENJAMIN A. WEAVER⁵

Draft version February 17, 2012

ABSTRACT

We present a hierarchical Bayesian determination of the velocity-dispersion function of approximately 430,000 massive luminous red galaxies (LRGs) observed at relatively low spectroscopic signal-to-noise ratio (SNR ~ 3 –5 per 69 km s^{-1}) by the Baryon Oscillation Spectroscopic Survey (BOSS) of the Sloan Digital Sky Survey III (SDSS-III). We marginalize over spectroscopic redshift errors, and use the full velocity-dispersion likelihood function for each galaxy to make a self-consistent determination of the velocity-dispersion distribution parameters as a function of absolute magnitude and redshift, correcting as well for the effects of broadband magnitude errors on our binning. Parameterizing the distribution at each point in the luminosity–redshift plane with a log-normal form, we detect significant evolution in the width of the distribution toward higher intrinsic scatter at higher redshifts. Using a subset of deep re-observations of BOSS galaxies, we demonstrate that our distribution-parameter estimates are unbiased regardless of spectroscopic SNR. We also show through simulation that our method introduces no systematic parameter bias with redshift. We highlight the advantage of the hierarchical Bayesian method over frequentist “stacking” of spectra, and illustrate how our measured distribution parameters can be adopted as informative priors for velocity-dispersion measurements from individual noisy spectra.

Subject headings: galaxies: evolution, kinematics and dynamics—methods: statistical—techniques: spectroscopic

1. INTRODUCTION

Massive elliptical galaxies (EGs; Hubble 1936) are one of the most important classes of astrophysical objects for galaxy evolution and cosmology. They represent the end stage of hierarchical galaxy-formation processes (e.g., Kauffmann et al. 1993; Baugh et al. 1996), and therefore their properties and scaling relations represent a key test for theories of galaxy formation within a cosmological context. In addition, since they are the most luminous and highly clustered galaxies, they serve as ideal cosmological tracers of clusters and large-scale structure (e.g., Eisenstein et al. 2005).

To a first approximation, EGs are “pressure-supported” rather than rotationally supported (e.g., Bertola & Capaccioli 1975; Illingworth 1977; Binney 1978), with their stellar motions characterized by a velocity dispersion σ . Among the many observational parameters of massive elliptical galaxies, σ is unique in its direct sensitivity to the depth of the galaxy’s gravitational potential (and therefore to its mass), and in its relatively weak dependence on observational aperture. In combination with galaxy sizes (i.e., half-light radii), velocity dispersions can be used to determine

“dynamical masses” that are independent of stellar-population assumptions (e.g., Padmanabhan et al. 2004; Bolton et al. 2008b). Dynamical masses can then in turn be used to trace the evolution of EGs at fixed mass (e.g., van der Marel & van Dokkum 2007; van der Wel et al. 2008; Cappellari et al. 2009), indicating a nuanced dynamical history despite generally passive star-formation histories at $z < 1$ (e.g., Thomas et al. 2005; Cool et al. 2008). Stellar velocity dispersion is also the most important single predictor of strong gravitational lensing cross sections (e.g., Turner et al. 1984; Bolton et al. 2008a), and can be used in combination with strong lensing observations to constrain the central mass-density structure of elliptical galaxies at cosmological distances (e.g., Koopmans & Treu 2002; Treu & Koopmans 2004; Koopmans et al. 2006). Stellar velocity dispersions are tied to nearly all other properties of EGs through multiple empirical scaling relations. Faber & Jackson (1976) found a correlation between luminosities of early-type galaxies and their velocity dispersions σ known as the Faber-Jackson Relation (FJR). The relation of Kormendy (1977) ties the surface brightness $\langle I \rangle_e$ with the effective radius R_e . Both the FJR and Kormendy relations can be viewed as projections of the “fundamental plane” (FP, e.g., Djorgovski & Davis 1987; Dressler et al. 1987; Bernardi et al. 2003c) within the space spanned by $\log_{10} R_e$, $\langle I \rangle_e$ and $\log_{10} \sigma$. Furthermore, central black hole mass has been found to be correlated with the velocity dispersion of the bulge via the $M_{\text{BH}} - \sigma$ relation (e.g., Ferrarese & Merritt 2000; Gebhardt et al. 2000; Kormendy & Bender 2009). Together, these relations provide multiple constraints on the structure, formation, and evolution of EGs.

¹ Department of Physics and Astronomy, University of Utah, 115 South 1400 East, Salt Lake City, UT 84112, USA (yiping.shu@utah.edu, bolton@astro.utah.edu)

² Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

³ Department of Astronomy, Yale University, New Haven, CT 06520 USA

⁴ Apache Point Observatory, Apache Point Road, P.O. Box 59, Sunspot, NM 88349, USA

⁵ Center for Cosmology and Particle Physics, New York University, NY 10003, USA

Although velocity dispersion plays a starring role in the study of EGs, it is an “expensive” observable that must be measured spectroscopically. Hence, large samples of galaxies with well-measured velocity dispersions across cosmic time are largely unavailable. Measurements of σ are made by quantifying the line-of-sight Doppler broadening of absorption lines relative to a set of template stellar spectra, either via the Fourier method (e.g., Sargent et al. 1977; Tonry & Davis 1979) or the direct-fitting method (e.g., Burbidge et al. 1961; Rix & White 1992). Both methods rely on the quality of galaxy spectra: for spectra of low signal-to-noise ratio (SNR), uncertainties in the measured stellar velocity dispersion can be very large and significantly non-Gaussian. This aspect is of particular concern for galaxies at cosmological distance (faint even if luminous), which can only be measured at high SNR through substantial investment of spectroscopic observing time and aperture.

In this paper, we introduce a hierarchical Bayesian statistical method to measure the parameters of the distribution of stellar velocity dispersions within a population of galaxies that has been observed with relatively low spectroscopic signal-to-noise ratio. We apply the method to approximately 430,000 luminous red galaxy (LRG) targets from the Baryon Oscillation Spectroscopic Survey (BOSS: Schlegel et al. 2009), one of four survey projects within the Sloan Digital Sky Survey III (SDSS-III: Eisenstein et al. 2011). We quantify the evolution of the velocity-dispersion function of BOSS galaxies, and detect significant evolution in the intrinsic population RMS of $\log_{10} \sigma$ at fixed absolute magnitude since $z \approx 0.8$.

This paper is organized as follows. In Section 2, we describe the sample selection and the method for velocity dispersion measurement. Section 3 presents our statistical method for the measurement of the distribution of stellar velocity dispersions within a population of galaxies, including a verification using high-SNR re-observations of a sub-sample of galaxies and a test for redshift-dependent systematic biases. Section 4 presents the results of our application of this method to the BOSS sample, showing the evolution of the velocity-dispersion function at fixed magnitude. Discussion and conclusions are presented in Section 5. Throughout the paper, we assume a standard general-relativistic cosmology with $\Omega_m = 0.3$, $\Omega_\Lambda = 0.7$ and $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

2. SPECTROSCOPIC DATA

2.1. Sample Selection

We use spectroscopic data obtained by the BOSS project via the 2.5-m SDSS telescope located at Apache Point Observatory in Sunspot, New Mexico (Gunn et al. 2006). The primary science goal of BOSS is the detection of the baryon acoustic feature in the two-point correlation function of galaxies (and quasar absorption systems), from which to constrain the distance–redshift relation and the nature of dark energy. BOSS also offers a unique resource for the study of the properties and evolution of massive galaxies. The BOSS footprint covers approximately $10,000 \text{ deg}^2$ in five imaging filters (*ugriz*, Fukugita et al. 1996), and will by 2014 obtain spectra of about 1.5 million LRGs out to redshift $z \approx 0.8$. Note that the majority of the BOSS LRG targets are massive EGs, although there is a small fraction of late-type galaxies as

well as unresolved multiples, particularly at the higher redshift end (Masters et al. 2011).

The BOSS spectra are broadly comparable to SDSS-I (York et al. 2000) spectra in resolution ($R \approx 2000$), and cover a wavelength range from 3,600Å to 10,000Å. The primary design goal of BOSS is to measure as many redshifts as efficiently as possible, in order to map the largest possible volume of the universe. Consequently, the SNR of the galaxy spectra is significantly lower than in SDSS-I, with typical SNR values of 3 to 5 per 69 km s^{-1} (re-binned) pixel, as compared with $\gtrsim 10$ per pixel in SDSS-I. Thus although the BOSS spectroscopic database is by far the largest available for the study of massive galaxies, the individual spectra are well below the SNR threshold of about 10 per Å generally regarded as a minimum for acceptable velocity-dispersion measurement on a galaxy-by-galaxy basis. Motivated by this context, we develop the Bayesian analysis method presented below.

Spectroscopic calibration, extraction, classification, and redshift measurement of all BOSS galaxy spectra are carried out using the `idlSpec2d` software (see, e.g., Aihara et al. 2011), written originally for SDSS-I and recently updated to handle the data format and noise regime of BOSS. In selecting our analysis sample, we make the following cuts based upon the redshift pipeline output:

- We use only the best spectroscopic observation of any given galaxy target as some objects are observed more than once (`SPECPRIMARY=1` according to SDSS terminology).
- We use only objects that were both targeted as galaxies and spectroscopically confirmed as galaxies
- We require a confident redshift measurement with no warning flags (`ZWARNING = 0` according to SDSS terminology).⁶

These cuts return approximately 430,000 galaxies from the first 1.5 years of BOSS spectroscopic observations, with redshifts ranging from zero to 1, but concentrated primarily over the interval $0.2 < z < 0.8$.

For all selected galaxies, we use the measured spectroscopic redshifts and SDSS broadband imaging colors to compute absolute rest-frame *V*-band magnitudes and associated uncertainties via the `sdss2bessell` routine implemented in the `kcorrect` software of Blanton & Roweis (2007).

The details of the BOSS galaxy target selection, and the corresponding incompletenesses, are the subject of a separate paper (Padmanabhan et al., in preparation). Here we summarize the target selection cuts for the two main galaxy target classes that we focus upon in our current study. The first is the CMASS sample (for “constant mass”), which is selected photometrically to deliver LRGs of approximately constant stellar mass over the redshift interval $0.3 < z < 0.8$, and which constitutes approximately 76% of the galaxies selected above. The second sample, LOZ, is selected to target LRGs at lower redshifts, and represents the remaining 24% of

⁶ For BOSS galaxies, the specifically relevant flag is `ZWARNING_NOQSO = 0` (Bolton et al., in prep.)




FIG. 1.— Distribution of galaxies for our sample, along with histograms of redshift z and V-band absolute magnitudes M_V . LOZ galaxies (black) and CMASS galaxies (gray) are plotted separately. For both the samples, contours are drawn at constant number density in the z - M_V plane, enclosing 50%, 90%, and 99% of the sample.

the selected galaxies. Defining the following quantities (Eisenstein et al. 2001; Cannon et al. 2006):

$$c_{\parallel} = 0.7(g - r) + 1.2[(r - i) - 0.18] \quad (1)$$

$$c_{\perp} = (r - i) - (g - r)/4.0 - 0.18 \quad (2)$$

$$d_{\perp} = (r - i) - (g - r)/8.0 \quad (3)$$

$$\text{ifiber2} = i\text{-band fiber magnitude for } 2'' \text{ fibers,} \quad (4)$$

the CMASS sample is defined by the photometric cuts:

$$17.5 < i < 19.9 \quad (5)$$

$$r - i < 2 \quad (6)$$

$$d_{\perp} > 0.55 \quad (7)$$

$$\text{ifiber2} < 21.7 \quad (8)$$

$$i < 19.86 + 1.60(d_{\perp} - 0.80) \quad (9)$$

as well as a cut to exclude galaxies with major-axis half-light radii greater than $8''$. Equations (5) and (7) aim to select galaxies between redshifts $z \sim 0.4$ – 0.8 , while Equation (9) attempts to impose a cut at constant stellar mass across this redshift range. The LOZ sample is defined by the cuts:

$$r < 13.5 + c_{\parallel}/0.3 \quad (10)$$

$$|c_{\perp}| < 0.2 \quad (11)$$

$$16 < r < 19.6. \quad (12)$$

Equation (10) sets up a magnitude threshold as a function of redshift and Equation (11) picks out low-redshift galaxies specifically.

The redshift–absolute magnitude distributions of these two BOSS galaxy samples, with associated 1D projections, are plotted in Figure 1. In the following analysis, we will treat the two populations separately, since the combined sample does not define a simple locus in luminosity–redshift space, with LOZ galaxies being of generally higher luminosity over the redshift range where the two samples overlap.

2.2. Velocity Dispersion Extraction

Our strategy for extracting velocity-dispersion information is to make use of the full velocity-dispersion likelihood function for each galaxy spectrum. To do this, we make use of the IDL routine `vdispfit` within the `idlpec2d` product of spectroscopic analysis software. This software has been extensively tested in the SDSS-I, and has been upgraded for the analysis of BOSS data. Velocity dispersions measured with this software have been the basis for multiple studies of the dynamics of EGs (e.g., Bernardi et al. 2003a,b,c; Sheth et al. 2003; Padmanabhan et al. 2004; Koopmans et al. 2006). To summarize briefly: `vdispfit` uses a set of stellar eigenspectra derived from a principal-component analysis (PCA) decomposition of the ELODIE stellar spectrum library (Prugniel & Soubiran 2001). The eigenspectra are convolved and binned to the resolution and sampling of the BOSS spectra, then broadened by Gaussian kernels of different trial velocity dispersions. The broadened templates are then shifted to the redshift of the galaxy under consideration. After masking out regions containing common emission lines, a linear least-squares fit is performed to obtain a best-fit model spectrum at each trial velocity dispersion. The resulting curve of χ^2 as a function of trial velocity dispersion encodes the likelihood function of velocity dispersion given the data. For measurements from high signal-to-noise spectra, the position of the minimum χ^2 is adopted as the maximum-likelihood estimate of the galaxy’s velocity dispersion. Below, rather than adopt these estimates, we will work with the full likelihood function.

In this procedure, we must choose the number of stellar eigenspectra to use in forming the template basis. The pipeline analysis of SDSS-I data used the first 24 PCA modes. For the much lower signal-to-noise BOSS data, an acceptable χ^2 can be obtained using only the first 5 PCA modes, and hence we restrict our basis to this smaller number of eigenspectra so as to avoid fitting noise fluctuations.

As described above, before being fit to the galaxy spectra, the stellar eigenspectra are shifted by the appropriate galaxy redshifts. If the redshifts have non-negligible errors, the corresponding offsets can introduce a bias into the measured velocity dispersion. Although the BOSS spectra provide redshifts with a precision well in excess of what is required for large-scale structure studies and absolute-magnitude determinations, their errors can be non-negligible on the scale of internal galaxy velocity dispersions. Therefore, we implement a marginalization over redshift errors in our analysis. Specifically, we modify the `vdispfit` routine to take a radial velocity-marginalization range Δz (expressed in constant-velocity pixels) and the redshift error δz (the $\pm 68\%$ confidence interval as estimated by the `idlpec2d` pipeline) as arguments. Then we calculate $\chi^2(\sigma, z)$ for a set of trial redshifts in the range $z \pm \Delta z$ and define a new effective $\chi^2(\sigma)$ by integrating over z as

$$\chi^2(\sigma) = -2 \ln \left(\int_{z-\Delta z}^{z+\Delta z} dz \exp \left[-\frac{\chi^2(\sigma, z)}{2} \right] p(z) \right), \quad (13)$$

We assume a Gaussian probability distribution for z

given by

$$p(z) \propto \exp \left[-\frac{(z - z_{\text{best}})^2}{2\delta z^2} \right] \quad (14)$$

where z_{best} is the best-estimate redshift from the BOSS spectroscopic pipeline. The choice of a Gaussian prior is made because the galaxy redshifts have been determined using absorption and emission-line information over the full optical range of the BOSS spectrograph, whereas the velocity-dispersion fitting is done only over the 4100–6800 Å rest-frame range covered by the ELODIE spectra, while also masking the wavelength positions of common emission lines. We also explored the use of a flat prior to completely marginalize over redshift in the velocity dispersion analysis, and found only a negligible change (at most a few percent) in the derived relations. For most galaxies, the effect of this marginalization on the χ^2 curve is insignificant, but since we wish to avoid introducing any spurious broadening into our population analysis, we apply the procedure to all spectra.

In this work, we do not make any aperture correction for velocity dispersions, although the angular BOSS fiber radius of 1'' subtends a different physical length scale as a function of redshift. Since aperture velocity dispersions are seen in the local universe to depend on aperture radius only to a weak power of approximately 0.04 to 0.06 (e.g., Jørgensen et al. 1995; Mehlert et al. 2003; Cappellari et al. 2006), this effect should be relatively insignificant. For example, taking a redshift range spanning the majority of our CMASS sample, the angular size of a fixed physical length at $z = 0.8$ is about 72% of its angular size at $z = 0.4$. Assuming the velocity dispersion within an aperture decreases as the aperture to the power -0.05 (a representative compromise value between the previous three references), this would correspond to a systematic change in measured velocity dispersion of about 1.7%, which is well below the level of other uncertainties in our analysis. In addition, the typical atmospheric seeing of approximately 1.8'' delivered to the BOSS spectroscopic focal plane will dilute the significance of the varying projected fiber scale. Essentially, BOSS velocity dispersions will represent a fair luminosity-weighted average value over the half-light radius of most target galaxies, which have half-light radii on the order of 1''.

3. STATISTICAL POPULATION ANALYSIS FORMALISM

The results of Bernardi et al. (2003b) suggest that the distribution of velocity dispersions for early-type galaxies at fixed luminosity can be well approximated by a log-normal function. Motivated by this, we will assume a Gaussian distribution in $\log_{10} \sigma$ with mean m and intrinsic scatter s :

$$p(\log_{10} \sigma | m, s) = \frac{1}{\sqrt{2\pi}s} \exp \left[-\frac{(\log_{10} \sigma - m)^2}{2s^2} \right] \quad (15)$$

We will treat m and s as functions of redshift and absolute magnitude, although we will suppress this dependence in our notation for convenience. Compared to the SDSS-I studies by Bernardi et al. (2003b) and Sheth et al. (2003), we have a much larger sample with greater redshift coverage, so we may investigate the evolution of both the mean and intrinsic scatter of $\log_{10} \sigma$




FIG. 2.— Results for stacking of spectra within a single redshift-luminosity bin. The top two panels show typical individual spectra, while the bottom panel shows the high-SNR stacked spectrum for that bin, resulting from averaging the spectra of ~ 200 galaxies.

with redshift and luminosity as encoded by these two distribution parameters (See also Bezanson et al. 2011 for a complementary analysis in terms of photometric velocity-dispersion proxies). Our strategy will be to analyze samples binned by an interval of 0.04 in redshift z , and by 0.1 in absolute magnitude M_V .

3.1. Frequentist Approach

As mentioned above, the SNR of BOSS galaxy spectroscopy is typically rather low, especially at the high-redshift end of the survey. Therefore, point estimation of the velocity dispersion of individual galaxies is of questionable reliability. Hence, we resort to analyzing the data by binning galaxies in the z - M_V plane, requiring at least 100 galaxies in every single bin. The most obvious first approach to determining the mean velocity dispersion in these bins is to remove the small relative redshift differences within the bin, stack all the spectra directly, and analyze the resulting high-SNR combination (see Figure 2). Although we do not adopt this method for our ultimate determinations of m and s , it is instructive to consider how such an approach relates to these parameters.

While a velocity dispersion can be measured at high SNR from the stacked spectrum, the measured value bears a non-trivial relation to the parameters m and s , which we now derive. Assuming equal luminosities within the bin (which basically holds by construction due to binning in absolute magnitude), what we measure from the stack σ_{stack}^2 is the population-weighted expectation value of σ^2 , i.e.

$$\sigma_{\text{stack}}^2 = \langle \sigma^2 \rangle = \int \sigma^2 p(\log_{10} \sigma | m, s) d \log_{10} \sigma. \quad (16)$$

The variance of σ_{stack}^2 is given by

$$\text{Var}(\sigma_{\text{stack}}^2) = \frac{1}{N} \text{Var}(\sigma^2) = \frac{1}{N} (\langle \sigma^4 \rangle - \langle \sigma^2 \rangle^2) \quad (17)$$

with N being the number of galaxies in the bin.

Making use of the following relation, which can be de-

rived for our log-normal form of Equation (15):

$$\begin{aligned} \langle \sigma^n \rangle &= \int \sigma^n p(\log_{10} \sigma | m, s) d \log_{10} \sigma \\ &= 10^{[n m + n^2 \ln(10) s^2 / 2]}, \end{aligned} \quad (18)$$

we find that

$$\sigma_{\text{stack}}^2 = 10^{[2 m + 2 \ln(10) s^2]} \quad (19)$$

$$\text{Var}(\sigma_{\text{stack}}^2) = \frac{(\sigma_{\text{stack}}^2)^2}{N} [10^{4 \ln(10) s^2} - 1]. \quad (20)$$

Thus we see that the velocity dispersion measured from the stacked spectrum is not given by the mean $\log\sigma$ value alone, but rather includes a contribution from the width of the population distribution as well. In principle, if a good estimator of $\text{Var}(\sigma_{\text{stack}}^2)$ can be obtained, the system can be closed and solved for m and s independently. Although we attempted to estimate $\text{Var}(\sigma_{\text{stack}}^2)$ via bootstrap resampling within each bin, we found the treatment of observational errors and varying signal-to-noise ratio among the spectra to be intractable within this framework. Rather than working further from measurements of stacked spectra, we proceed to the hierarchical Bayesian method described in the following section.

3.2. Hierarchical Bayesian Approach

To constrain the distribution parameters m and s within each redshift–magnitude bin, we consider the following expansion of the likelihood function $\mathcal{L}(m, s | \{\vec{d}\})$ in the bin:

$$\begin{aligned} \mathcal{L}(m, s | \{\vec{d}\}) &= p(\{\vec{d}\} | m, s) \\ &= \prod_i p(\vec{d}_i | m, s) \\ &= \prod_i \int p(\vec{d}_i | \log_{10} \sigma) p(\log_{10} \sigma | m, s) d \log_{10} \sigma \end{aligned} \quad (21)$$

Here $\{\vec{d}\}$ is the set of all spectra in the bin, with each element \vec{d}_i representing the spectrum of the i^{th} galaxy. The expression $p(\vec{d}_i | \log_{10} \sigma)$ is related to the $\chi^2(\log_{10} \sigma)$ function by

$$p(\vec{d}_i | \log_{10} \sigma) \propto \exp \left[-\frac{\chi_i^2(\log_{10} \sigma)}{2} \right], \quad (22)$$

and $p(\log_{10} \sigma | m, s)$ is given by Equation (15). Translating into Bayesian terms, we have a posterior probability for m and s given by

$$p(m, s | \{\vec{d}\}) \propto p(\{\vec{d}\} | m, s) p(m, s) \quad (23)$$

with $p(m, s)$ being the prior probability distribution for m and s . For simplicity, we assume a uniform prior on m and s over a reasonable range. In actuality, we find that the likelihood is quite strongly peaked in each bin, so the exact nature and range of the prior are insignificant.

3.3. Verification

To verify the correct functioning of our Bayesian approach, we make use of data from BOSS plate 3851. Due to a CCD failure on one of the two BOSS spectrographs




FIG. 3.— Engineered sub-sample Gaussian histogram in $\log_{10} \sigma$ constructed using measurements from 7-hour BOSS observations (black), with histogram of same sub-sample using velocity dispersions measured from 1.75-hour observations (gray). The two histograms have been given a slight relative horizontal offset, for display purposes. The vertical dashed line indicates the mean $\log_{10} \sigma$ value of 2.33 for the sub-sample. Note the relative broadening of the 1.75-hour histogram due to the effects of observational error.

that temporarily suspended normal survey operations, 500 of the 1000 targets on this plate were plugged and observed for a total integration time of 7 hours (28 exposures of 15 minutes each) over the course of several nights ending on 2010 April 12, denoted within the SDSS-III database by the modified Julian date (MJD) of 55298. Subsequent to the replacement of the failed CCD, the entire plate was re-plugged and observed for a more typical BOSS integration time of 1.75 hours total on MJD 55302. The set of re-observed targets allows us to compare BOSS galaxy spectra of standard survey depth with spectra of the same objects at essentially double the nominal survey SNR. We use these repeat observations to verify that our method (1) does not have a signal-to-noise ratio dependent bias in the estimation of velocity-dispersion distribution parameters, and (2) reproduces the known distribution of velocity dispersions within a controlled sample, as measured from the high-SNR set of spectra.

Between the deep and shallow re-observations, there are 308 galaxies which have equal redshifts (within $\Delta z = \pm 0.005$) and classifications for both observation dates. Since the sample is heterogeneous in magnitude and redshift, we select a sub-sample with a controlled distribution in velocity dispersion. We restrict our attention to galaxies that have their individual velocity dispersions measured at SNR of 10 or more from the 7-hour observations, and that have redshifts between 0.4 and 0.6. We then select a random sub-sample of 125 galaxies from this set so as to have a Gaussian histogram in $\log_{10} \sigma$ with a mean of $m = 2.33$ and an intrinsic RMS scatter of $s = 0.07$. The histogram of this sub-sample, along with the histogram of the same sample as constructed from galaxy-by-galaxy measurements using the 1.75-hour observations, is shown in Figure 3.

The frequentist formulas given by Equations (19) and (20) do not account for observational error, and hence we do not use them to solve for m and s estimates for our relatively low signal-to-noise BOSS survey data. However, our subsample of high signal-to-noise 7-hour observations allows us to test them, which we do before proceeding to the verification of our Bayesian analysis framework. First, we use Equation (19) with a mean of




FIG. 4.— Credible-region contours of constant posterior probability density for m and s parameters measured from the engineered test sub-sample of galaxies observed with both 7-hour integrations (black) and 1.75-hour integrations (gray). The symbol is the location of the parameters chosen for the construction of the test sub-sample. The offset in s between the contours and the symbol is a result of the proper deconvolution of observational uncertainty that is implemented by the Bayesian method.

2.33 and an intrinsic scatter of 0.07 to predict a value of $\sigma_{\text{stack}} = 219 \text{ km s}^{-1}$, which is in very good agreement with the result of $(222 \pm 12) \text{ km s}^{-1}$ that we obtain by fitting the stacked spectrum of this set of 125 galaxies directly. Similarly, we predict $[\text{Var}(\sigma_{\text{stack}}^2)]^{1/4} = 40 \text{ km s}^{-1}$ from Equation (20), which is in reasonable agreement with the value of 46 km s^{-1} obtained through a bootstrap resampling process. In both cases the agreement is not exact because there is still some observational error even in the 7-hour data, but as mentioned above, we will pass to the Bayesian framework to quantify these effects.

We next carry out the estimation of the m and s parameters of the selected sub-sample of objects, using the Bayesian approach described above, for both the 7-hour and the 1.75-hour data sets. Figure 4 shows the resulting posterior probability density for these parameters as estimated from both data sets. As expected, we see that the posterior PDF is tighter for the 7-hour data. More importantly, we see no significant bias in the posterior PDF between the low-SNR and high-SNR data sets. This is especially significant for the estimation of the s parameter: if we were not handling our observational uncertainties correctly, we might expect to infer a broader intrinsic distribution (higher s value) from the noisier data, but this not the case. We also see that the parameters used to engineer the subsample are recovered with no significant bias in m . We see a slight offset of the 7-hour maximum posterior s value from the input value used to engineer the sample. This is in the direction and of the size to expected given the observational error of the 7-hour individual-spectrum velocity dispersion measurements, which have an RMS signal-to-noise of about 17. This corresponds to an observational broadening of about 0.025 dex in the engineered histogram of Figure 3, which is deconvolved by the Bayesian parameter estimation procedure to give the lower recovered s value seen in Figure 4.

Another concern is that there might be a systematic bias with redshift, since the spectral regions used by `vdispfit` in fitting for velocity dispersions (rest frame wavelength range from $4,100\text{\AA}$ to $6,800\text{\AA}$) move to the

redder and noisier parts of the spectrum as the redshift gets higher. In order to test this, we construct another controlled sub-sample with 152 galaxies of redshift $z < 0.2$ and very high SNR. Then we take the best-fit template combination models of those 152 galaxies returned by `vdispfit` and redshift them to progressively higher redshift bins, giving them a uniform random distribution over a bin width of $\Delta z = 0.04$ in each case (to match our actual binning). At each new redshift, the model spectra are added to sky-subtracted BOSS sky fibers to simulate realistic survey noise, and scaled individually in flux to give a typical median SNR at that redshift bin. We then analyse the simulated redshifted samples with our Bayesian method to estimate the posterior PDF of m and s . The results are shown in Figure 5 and Figure 6, for 5 separate redshift bins. We see that the recovered parameters are consistent within observational error across all redshifts, with no apparent redshift-dependent bias.




FIG. 5.— Contours of constant posterior probability density (68%, 95%, and 99%) for m and s parameters obtained from a controlled sub-sample of 152 galaxies in 5 different redshift bins with gradually reduced SNRs.




FIG. 6.— The best estimated m (Diamond) and s (Triangle) values for the controlled sub-sample of 152 galaxies at 5 different redshift bins.

Finally, to rule out any significant dependence of our measurement on airmass and fiber position within the BOSS spectrographs, we make use of data from plates

3615, 3647, 4238 and 4239. These four plates cover roughly the same area of sky, but with different plate drillings that place the same objects in very different fibers within the spectrograph system. They were also observed over a range of different airmasses on multiple nights. From these plates, we construct several subsamples of spectra, all of which include the same galaxies, but are drawn from different plates and/or observations. As with the previous tests, we recover consistent estimates of m and s from the analysis of all these samples.

Based on the above three tests, we conclude that our method recovers accurate estimates of the population velocity-dispersion distribution parameters.

3.4. Magnitude Error Correction

Our method of determining $p(\vec{d}_i | \log_{10} \sigma)$ incorporates an explicit marginalization over redshift error, and propagates all observational uncertainty in the velocity dispersion of a given galaxy. Our binning in redshift and absolute magnitude introduces additional error possibilities that we must account for. In the case of redshift, the errors are negligible relative to the bin width of $\Delta z = 0.04$, and are unlikely to contribute any artificial broadening to our determination of the redshift dependence of m and s . The absolute magnitude errors are, however, non-negligible in comparison to the magnitude bin width of $\Delta M_V = 0.1$, and thus we use the following technique to estimate and compensate for the broadening effect of the observational scattering of galaxies between absolute-magnitude bins (see Figure 1).

Suppose (m, s) are the true values within a bin, and (m_1, s_1) are the values that we determine in the presence of absolute-magnitude errors. We assume that

$$m_1 = m + \delta m \quad (24)$$

$$s_1^2 = s^2 + \delta s^2, \quad (25)$$

where δm and δs are the biases introduced by magnitude errors. To estimate and remove these biases, we add additional random errors to all our galaxy absolute magnitudes M_V to give

$$M'_V = M_V + \epsilon \delta M_V, \quad (26)$$

where ϵ is a normally distributed random number with mean 0 and standard deviation 1, and δM_V are the galaxy-by-galaxy absolute-magnitude errors estimated by `sdss2besse11` (propagated from SDSS *ugriz* apparent magnitude errors). We repeat our analysis, binning instead in M'_V , and denoting the new distribution parameter results by m_2 and s_2 . We assume these new determinations are related to m_1 and s_1 in the same way as m_1 and s_1 are related to m and s , which implies that

$$m_2 = m + 2 \delta m \quad (27)$$

$$s_2^2 = s^2 + 2 \delta s^2 \quad (28)$$

Thus the biases due to absolute magnitude errors δm and δs can be removed to yield

$$m = 2 m_1 - m_2 \quad (29)$$

$$s = \sqrt{2 s_1^2 - s_2^2} \quad (30)$$

In practice, we find typical values for δm of 0.01, and for δs of 0.04.

4. RESULTS: EVOLUTION OF THE VELOCITY-DISPERSION FUNCTION

In this section we present the results of the application of our hierarchical Bayesian velocity-dispersion distribution measurement technique to the approximately 103,000 galaxies in our LOZ sample and 330,000 galaxies in our CMASS sample.

4.1. LOZ Sample

The LOZ sample extends to $z \approx 0.5$. The 2D contour plots of m and s (Figure 7) and scatter plots in different redshift bins (Figure 8) show that the mean m is strongly correlated with absolute magnitude, while the intrinsic scatter s shows no significant variation. Tracks of constant stellar mass assuming the LRG stellar population model of Maraston et al. (2009) have also been over-plotted in Figure 7, and used to convert from an absolute-magnitude to a stellar-mass baseline in Figure 8. Galaxies in the LOZ sample have estimated stellar masses between approximately $10^{11} M_\odot$ and $10^{12} M_\odot$.

To quantify the variation of the m and s parameters with redshift and absolute magnitude, we consider a simple model specified by:

$$m^0 = A_m^0 M_V + B_m^0 \log_{10}(1+z) + C_m^0 \quad (31)$$

$$s^0 = A_s^0 M_V + B_s^0 \log_{10}(1+z) + C_s^0, \quad (32)$$

with the “0” superscript denoting the LOZ sample specifically. Performing a linear least squares fit to the individual bin data points, we obtain

$$\begin{aligned} A_m^0 &= -0.0880 \pm 0.0012 & A_s^0 &= 0.006 \pm 0.002 \\ B_m^0 &= -0.087 \pm 0.018 & B_s^0 &= -0.08 \pm 0.03 \\ C_m^0 &= 0.37 \pm 0.02 & C_s^0 &= 0.20 \pm 0.04 \end{aligned} \quad (33)$$

We can translate the resulting scaling into the standard form for the FJR, with luminosity $L \propto \sigma^x$ by recognizing that $x = -0.4/A_m^0$. The resulting value of $x = 4.55 \pm 0.06$ is in reasonable agreement with the canonical local-universe value of $x = 4$. Thus, BOSS low- z LRGs define an FJR whose slope and scatter has little dependence on redshift and luminosity; there is correspondingly little evidence for dynamical evolution in this sample since roughly $z = 0.5$.

4.2. CMASS Sample

The CMASS galaxy sample extends from $z \approx 0.3$ to $z \approx 0.8$. The results of our σ distribution parameter measurements are shown in Figures 9 and 10, once again using tracks of constant stellar mass based on the Maraston et al. (2009) population model. Using the same model form as used for the LOZ sample above,

$$m^1 = A_m^1 M_V + B_m^1 \log_{10}(1+z) + C_m^1 \quad (34)$$

$$s^1 = A_s^1 M_V + B_s^1 \log_{10}(1+z) + C_s^1 \quad (35)$$

(with the “1” superscript denoting the CMASS sample specifically), and again doing a linear least-squares fit, we find that

$$\begin{aligned} A_m^1 &= -0.1128 \pm 0.0010 & A_s^1 &= 0.0263 \pm 0.0016 \\ B_m^1 &= -0.77 \pm 0.02 & B_s^1 &= 0.82 \pm 0.04 \\ C_m^1 &= -0.089 \pm 0.019 & C_s^1 &= 0.52 \pm 0.03 \end{aligned}$$

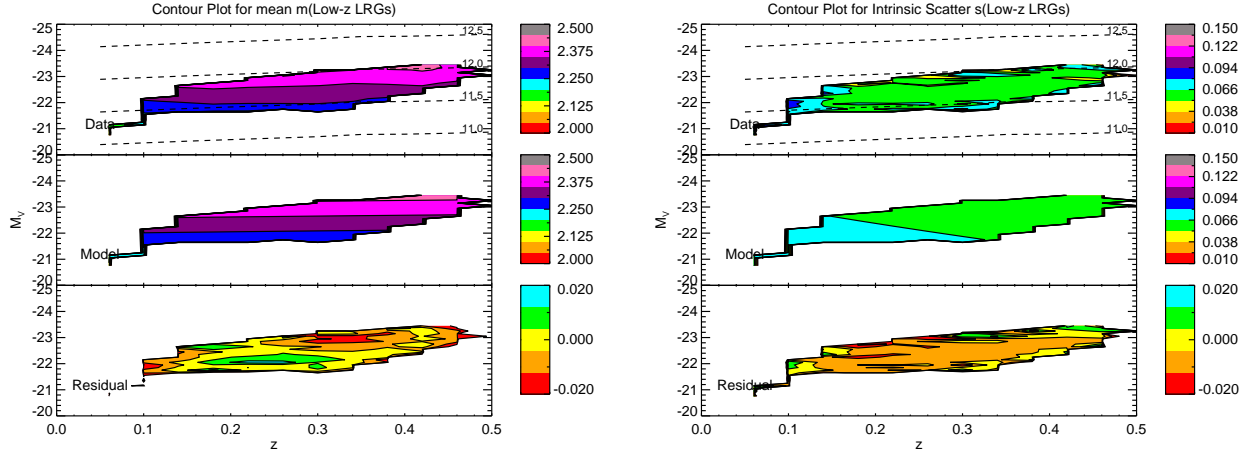


FIG. 7.— Contour plots of m & s for LOZ sample galaxies. Top panels show the map of maximum posterior probability over the range of the plane with bins containing at least 100 galaxies. Middle panels show low-order bivariate model fits to these maps constructed as described in the text, and residuals (top minus middle) are shown in the bottom panels. Dashed lines in the top panels show tracks of constant stellar mass from the LRG population model of Maraston et al. (2009).




FIG. 8.— Scatter plots of m & s versus M_V (top two panels) and $\log_{10}(M_*/M_\odot)$ (bottom two panels) for LOZ sample galaxies in different redshift ranges.

In the case of the CMASS sample, the FJR is still apparent, but the scaling exponent in $L \propto \sigma^x$ is now $x = 3.55 \pm 0.03$. This observation that the FJR becomes “shallower” at higher redshift can be interpreted in terms of mass-dependent star-formation history (e.g., Cowie et al. 1996; di Serego Alighieri et al. 2005), with less massive (lower σ) galaxies having undergone more recent star formation and thus fading more rapidly with cosmic time relative to more massive galaxies.

There is a clear evolution in the zero-point of the m

versus M_V relation (upper left panel in Figure 10) with redshift. This evolution is essentially eliminated in the lower left panel of Figure 10, which translates to a baseline of constant stellar mass. Hence the evolution of the m versus M_V relation in the CMASS sample is consistent with passive stellar evolution.

It can easily be seen from Figures 9 and 10 that s is no longer constant with redshift at fixed luminosity or stellar mass. The significance of this result is encapsulated in the non-zero value of $B_s^1 = 0.82 \pm 0.04$ given above. To

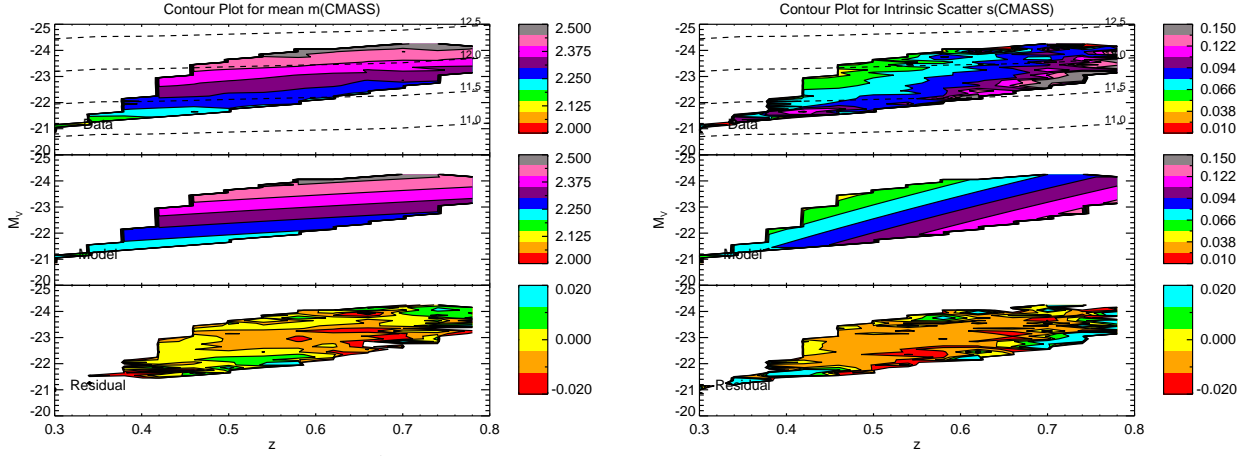


FIG. 9.— The same as Figure 7 but for CMASS galaxies.




FIG. 10.— The same as Figure 8 but for CMASS galaxies. (Note that the scales of these panels are expanded relative to Figure 8.) The increase of intrinsic scatter with redshift can be seen in the right-hand figure.

quantify this result in more detail, we fit the s versus M_V relation with a linear model at each redshift bin, and plot the zero-point of this relation as a function of redshift in Figure 11. We see that within the CMASS sample, the intrinsic width s of the velocity-dispersion function at fixed magnitude or stellar mass *decreases* with cosmic time (i.e., broader distribution width at higher redshift), especially at redshifts $z > 0.6$. This is consistent with our tentative detection of evolution in the FJR slope between the LOZ and CMASS samples, in the sense that a given range in luminosity encompasses a larger range of velocity dispersions at higher redshift, but the signal

is too large to be explained by this effect alone (since the FJR slope is not seen to evolve significantly *within* the CMASS sample alone). We are therefore seeing increased dynamical heterogeneity at fixed luminosity in the CMASS sample at higher redshifts.

We note that the apparent increase in the intrinsic σ distributions at high redshift cannot be explained in terms of surface-brightness selection effects. Through the FP or Kormendy relations, velocity dispersion at fixed luminosity is correlated with surface brightness. At the high-redshift end of the CMASS sample, we can expect a degree of incompleteness at both ends of the surface-




FIG. 11.— Variation of the intrinsic width s of the CMASS population distribution in $\log_{10} \sigma$ as a function of redshift.

brightness distribution. On the one hand, relatively low surface-brightness galaxies will have fainter magnitudes within the BOSS spectroscopic fiber, and will thus be less likely to be targeted, and less likely to have confident and correct spectroscopic redshift measurements even if targeted. On the other hand, relatively high surface-brightness galaxies (again, at fixed luminosity) run the risk of being unresolved in star–galaxy separation. Consequently, we might expect the distribution of velocity dispersion at fixed magnitude to be made more narrow at high redshift by these considerations, which goes in the opposite sense to the trend we detect.

4.3. Application to Individual Spectra

Our results characterize the dynamical properties of the population of LRGs targeted by BOSS. The parameters of the population can in turn be used to inform our estimates of the velocity dispersion values of individual noisy BOSS spectra. For this application, we want to use distribution parameters *uncorrected* for broadband magnitude errors, since these same errors will be present in the photometric data for the individual galaxies whose spectra we wish to analyze.

For low- z LRGs, without magnitude error correction, we have

$$m^0 = -0.0829M_V - 0.042 \log_{10}(1+z) + 0.48 \quad (36)$$

$$s^0 = 0.006M_V - 0.09 \log_{10}(1+z) + 0.22 \quad (37)$$

and for CMASS galaxies, without magnitude error correction, we have

$$m^1 = -0.0973M_V - 0.60 \log_{10}(1+z) + 0.23 \quad (38)$$

$$s^1 = 0.0240M_V + 0.76 \log_{10}(1+z) + 0.49 \quad (39)$$

We can then take the posterior probability $p(m, s | \{\vec{d}\})$ from the entire sample as a prior probability for the analysis of an individual galaxy spectrum. The posterior probability for $\log_{10} \sigma$ of the spectrum is then

$$p(\log_{10} \sigma | \vec{d}_i) \propto p(\vec{d}_i | \log_{10} \sigma) p(m, s | \{\vec{d}\}) . \quad (40)$$

Loosely speaking, if the observational error in the velocity dispersion measured from a single spectrum is comparable to the intrinsic width s of the particular population from which it is drawn, then the data and the prior will contribute equally to the determination of the posterior

PDF of $\log_{10} \sigma$. If the observational error is small, the effect of the prior will be correspondingly minor, while if the observational error is large, the posterior PDF will be determined primarily by the prior.

The application of this method can thus permit a more precise σ estimate for individual galaxies, by making use of the collective information about the population from which it was drawn. It is important however to note that if the spectra under consideration are somehow selected to be biased towards either higher or lower velocity dispersions, then the prior will pull them systematically towards the population mean, giving posterior PDFs that are biased relative to the true σ values. We must also be sure only to apply this method to subsamples of spectra that are much smaller than the population samples used to determine the distribution parameters.

5. DISCUSSION & CONCLUSION

In this paper, we have presented a new technique for estimating the velocity-dispersion function of LRGs from large numbers of low SNR spectra. This method incorporates the effects of observational uncertainties in spectroscopic redshift, velocity dispersion, and broadband magnitude. We have compared our method favorably to the more traditional approach of “stacking” multiple spectra; our new approach can perhaps be termed “Bayesian stacking”. We have also indicated how the results of our method can be used as informative priors to provide more precise estimates of the velocity dispersions of individual galaxies, provided that those galaxies are an unbiased selection from the parent distribution at their particular redshift and luminosity.

We have applied our technique to a sample of 430,000 galaxy spectra from the BOSS project of the SDSS-III, covering the redshift range from zero to unity, concentrated between approximately $z = 0.2$ and $z = 0.8$. For the higher-redshift CMASS target sample (approximately 76% of our galaxies), we detect a highly significant increase in the intrinsic width of the velocity-dispersion distribution at higher redshifts, indicative of greater galaxy diversity at fixed luminosity at earlier cosmic times. For the lower-redshift LOZ galaxy sample, we find little evolution in the velocity-dispersion distribution below $z \approx 0.5$. Although the CMASS and LOZ samples do not form a single uniform sample (LOZ galaxies being generally more luminous than CMASS targets over the range of redshift where the two overlap), our results suggest that dynamical evolution of massive LRGs is much more significant over the interval $0.5 < z < 1.0$ as compared to $0 < z < 0.5$.

Future applications of this method to the BOSS galaxy samples will focus on the effects of observational selection on the deduced population evolution. We also plan to divide our analysis further by rest-frame color, so as to differentiate between galaxies of different stellar population at a given redshift and magnitude. By making a more accurate division of the sample in terms of stellar mass and star-formation history, we hope to separate the signatures of dynamical and stellar-population evolutionary channels, and to thereby obtain a more detailed picture of LRG population evolution and a more powerful discriminant between theoretical scenarios. This approach can also determine whether the effect of increased population scatter in $\log_{10} \sigma$ at high redshift is due to greater

dynamical diversity, greater stellar-population diversity, or to some combination of the two effects.

Our measurements can also have important implications for the statistics of gravitational lensing, by constraining the total lensing cross-section in massive elliptical galaxies between redshift 0 and 1. Although a precise application to gravitational-lensing statistics must await a proper treatment of completeness, our current results can be combined with published luminosity functions (e.g., Cimatti et al. 2006; Cool et al. 2008) to place a lower limit on the integrated lensing cross section.

The application of hierarchical Bayesian methods such as the one presented here may hold the key to reconciling the tension between redshift surveys designed for constraining cosmological parameters and those designed for the study of galaxy evolution. The former goal generally dictates an SNR just sufficient to measure redshift for as many galaxies over as large a volume of the universe as possible, while the latter goal traditionally requires observations at high enough SNR to precisely constrain multiple physical parameters for each galaxy. However, if the ultimate goal of galaxy-evolution studies is to measure the distribution of physical parameters within a statistically significant sample of galaxies, then Bayesian methods can remove the need to measure those parameters precisely on a galaxy-by-galaxy basis. In fact, there may indeed be an objective galaxy-evolution case for trading fewer high-SNR spectra for more low-SNR

spectra, so as to reduce the effects of sample variance. If cosmological experimental designs can also accommodate the more permissive (e.g., magnitude-limited) targeting desired for galaxy population studies, then both goals may be well served by the same redshift survey.

The authors thank David W. Hogg and Glenn van de Ven for valuable discussion of this work.

Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The SDSS-III web site is <http://www.sdss3.org/>.

SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the University of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, University of Cambridge, University of Florida, the French Participation Group, the German Participation Group, the Instituto de Astrofísica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, New Mexico State University, New York University, Ohio State University, Pennsylvania State University, University of Portsmouth, Princeton University, the Spanish Participation Group, University of Tokyo, The University of Utah, Vanderbilt University, University of Virginia, University of Washington, and Yale University.

REFERENCES

- Aihara, H., et al. 2011, *ApJS*, 193, 29
 Baugh, C. M., Cole, S., & Frenk, C. S. 1996, *MNRAS*, 283, 1361
 Bernardi, M., et al. 2003a, *AJ*, 125, 1817
 —. 2003b, *AJ*, 125, 1849
 —. 2003c, *AJ*, 125, 1866
 Bertola, F., & Capaccioli, M. 1975, *ApJ*, 200, 439
 Bezanson, R., et al. 2011, *ApJ*, 737, L31
 Binney, J. 1978, *MNRAS*, 183, 501
 Blanton, M. R., & Roweis, S. 2007, *AJ*, 133, 734
 Bolton, A. S., et al. 2008a, *ApJ*, 682, 964
 —. 2008b, *ApJ*, 684, 248
 Burbidge, E. M., Burbidge, G. R., & Fish, R. A. 1961, *ApJ*, 133, 393
 Cannon, R., et al. 2006, *MNRAS*, 372, 425
 Cappellari, M., et al. 2006, *MNRAS*, 366, 1126
 —. 2009, *ApJ*, 704, L34
 Cimatti, A., Daddi, E., & Renzini, A. 2006, *A&A*, 453, L29
 Cool, R. J., et al. 2008, *ApJ*, 682, 919
 Cowie, L. L., Songaila, A., Hu, E. M., & Cohen, J. G. 1996, *AJ*, 112, 839
 di Serego Alighieri, S., et al. 2005, *A&A*, 442, 125
 Djorgovski, S., & Davis, M. 1987, *ApJ*, 313, 59
 Dressler, A., Lynden-Bell, D., Burstein, D., Davies, R. L., Faber, S. M., Terlevich, R., & Wegner, G. 1987, *ApJ*, 313, 42
 Eisenstein, D. J., et al. 2001, *AJ*, 122, 2267
 —. 2005, *ApJ*, 633, 560
 —. 2011, *ArXiv e-prints*
 Faber, S. M., & Jackson, R. E. 1976, *ApJ*, 204, 668
 Ferrarese, L., & Merritt, D. 2000, *ApJ*, 539, L9
 Fukugita, M., Ichikawa, T., Gunn, J. E., Doi, M., Shimasaku, K., & Schneider, D. P. 1996, *AJ*, 111, 1748
 Gebhardt, K., et al. 2000, *ApJ*, 539, L13
 Gunn, J. E., et al. 2006, *AJ*, 131, 2332
 Hubble, E. P. 1936, *Realm of the Nebulae* (New Haven: Yale University Press)
 Illingworth, G. 1977, *ApJ*, 218, L43
 Jørgensen, I., Franx, M., & Kjaergaard, P. 1995, *MNRAS*, 276, 1341
 Kauffmann, G., White, S. D. M., & Guiderdoni, B. 1993, *MNRAS*, 264, 201
 Koopmans, L. V. E., & Treu, T. 2002, *ApJ*, 568, L5
 Koopmans, L. V. E., Treu, T., Bolton, A. S., Burles, S., & Moustakas, L. A. 2006, *ApJ*, 649, 599
 Kormendy, J. 1977, *ApJ*, 218, 333
 Kormendy, J., & Bender, R. 2009, *ApJ*, 691, L142
 Maraston, C., Strömbäck, G., Thomas, D., Wake, D. A., & Nichol, R. C. 2009, *MNRAS*, 394, L107
 Masters, K. L., et al. 2011, *ArXiv e-prints*
 Mehlert, D., Thomas, D., Saglia, R. P., Bender, R., & Wegner, G. 2003, *A&A*, 407, 423
 Padmanabhan, N., et al. 2004, *New Astron.*, 9, 329
 Prugniel, P., & Soubiran, C. 2001, *A&A*, 369, 1048
 Rix, H.-W., & White, S. D. M. 1992, *MNRAS*, 254, 389
 Sargent, W. L. W., Schechter, P. L., Boksenberg, A., & Shortridge, K. 1977, *ApJ*, 212, 326
 Schlegel, D., White, M., & Eisenstein, D. 2009, in *Astronomy, Vol. 2010, astro2010: The Astronomy and Astrophysics Decadal Survey*, 314–+
 Sheth, R. K., et al. 2003, *ApJ*, 594, 225
 Thomas, D., Maraston, C., Bender, R., & Mendes de Oliveira, C. 2005, *ApJ*, 621, 673
 Tonry, J., & Davis, M. 1979, *AJ*, 84, 1511
 Treu, T., & Koopmans, L. V. E. 2004, *ApJ*, 611, 739
 Turner, E. L., Ostriker, J. P., & Gott, III, J. R. 1984, *ApJ*, 284, 1
 van der Marel, R. P., & van Dokkum, P. G. 2007, *ApJ*, 668, 756
 van der Wel, A., Holden, B. P., Zirm, A. W., Franx, M., Rettura, A., Illingworth, G. D., & Ford, H. C. 2008, *ApJ*, 688, 48
 York, D. G., et al. 2000, *AJ*, 120, 1579